

# Video and Image Analytics for Marine Environments (VIAME) An Open-Source, Do-It-Yourself AI Toolkit

Matt Dawkins (Speaker), Senior R&D Engineer, [matt.dawkins@kitware.com](mailto:matt.dawkins@kitware.com)

Other Developers: Aashish Chaudhary, Jon Crall, Matthew Woehlke, Bryon Lewis, Jacob Nesbitt, Brandon Davis, Neal Siekierski, Anthony Hoogs, Linus Sherrill, Matt Brown, Betsy McPhail, Kyle Edwards, Matt Leotta, Rusty Blue

NOAA: Benjamin Richards, Dvora Hart, George (Randy) Cutter, Elizabeth Clarke, Charles Thompson, Kresimir Williams, Bill Michaels, Erin Moreland, Katie Sweeney, Abigail Powell

CFF: Liese Siemann




Kitware

# What is VIAME?



- A do-it-yourself (DIY) AI toolkit which can be applied to multiple types of imagery or video
- Can be run by people with no programming or machine learning background in both web and desktop interfaces
- Released as fully open-source with a permissive license, see [viametoolkit.org](https://viametoolkit.org)
- Specializations to maritime processing such as motion fusion, stereo measurement, image enhancement, and object tracking which other software (e.g. Amazon SageMaker) lack

Object Detection



Object Tracking






Image Enhancement




Image Registration and Mosaicing




Video Search and Rapid Model Generation




Stereo Measurement



Calibration, 3D and Altitude Estimation









# Aerial Surveys (Manned Fixed-Wings, UAVs)




Source: AKFSC MML Stellar Sea Lion



Source: AKFSC MML Arctic Seal



Source: AKFSC MML Stellar Sea Lion



Source: NEFSC Gray Seal




Source: SWFSC



Source: NEFSC Harbor Seal




Source: SWFSC Penguin Aerial Data



Sources: University of Alaska Data  
Scene Segmentation

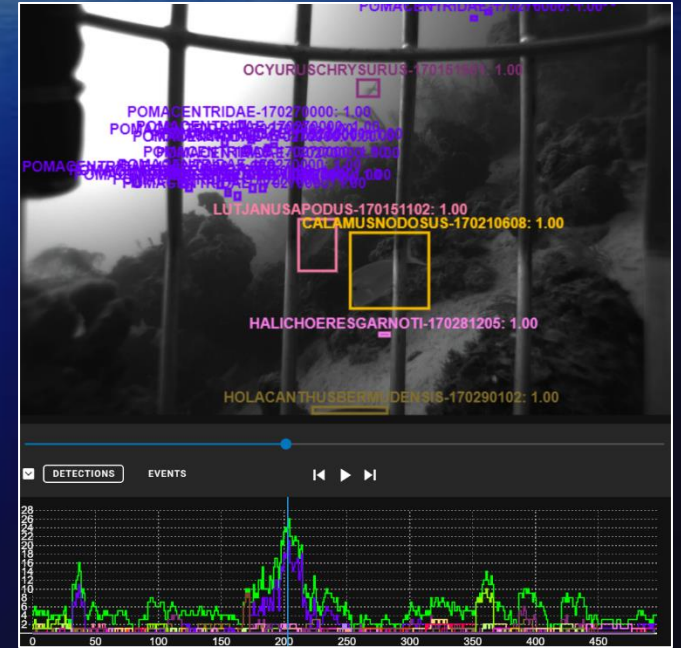
## Outward Facing Cameras (Underwater, Ship-Based)



**Source: PIFSC MOUSS**



**Source: SWFSC Penguin Cam**



**Source: SEFSC Quadcam**





**Source: PIFSC MOUSS**



**Source: PIFSC EM Data**



# Example Platforms

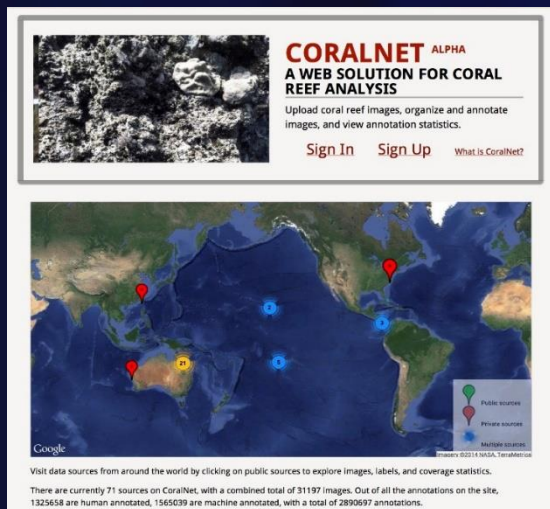


# NMFS Strategic Initiative on Automated Image Analysis


**Mission:** Develop guidelines, set priorities, and fund projects to develop broad-scale, standardized, and efficient automated analysis of still and video imagery for use in underwater stock assessment

**Funded VIAME and CoralNet from 2015 to present**

- Benjamin Richards (*chair*)  
NOAA Pacific Islands Fisheries Science Center
- Alexandra Branzan Albu  
University of Victoria
- Elizabeth Clarke  
NOAA Northwest Fisheries Science Center
- George “Randy” Cutter  
NOAA Southwest Fisheries Science Center
- Duane Edgington  
Monterey Bay Aquarium Research Institute
- Dvora Hart  
NOAA Northeast Fisheries Science Center
- Anthony Hoogs  
Kitware, Inc.
- David Kriegman  
University of California, San Diego
- Clay Kunz  
Google
- Michael Piacentino  
SRI International
- Lakshman Prasad  
Los Alamos National Laboratory
- Charles Thompson  
NOAA Southeast Fisheries Science Center
- Kresimir Williams  
NOAA Alaska Fisheries Science Center




<http://coralnet.ucsd.edu>, D. Kriegman, UCSD



Flask, Michael Piacentino, SRI



Misc. Analytics (e.g. LANL)




**2019 Department of  
Commerce Gold Medal  
Awarded to NOAA Members of  
AIASI for VIAME and CoralNet**

**Other Funded Initiatives**



# Kitware

- Collaborative software R&D: algorithms & applications, image & data analysis, support & training
- Best known for open source toolkits and applications
- 150+ employees:
  - $\frac{1}{3}$  masters
  - $\frac{1}{3}$  PhD
- Founded in 1998
- Offices in Albany, NY; Chapel Hill, NC; Santa Fe, NM; Minneapolis, MN; Arlington, VA; Lyon, France



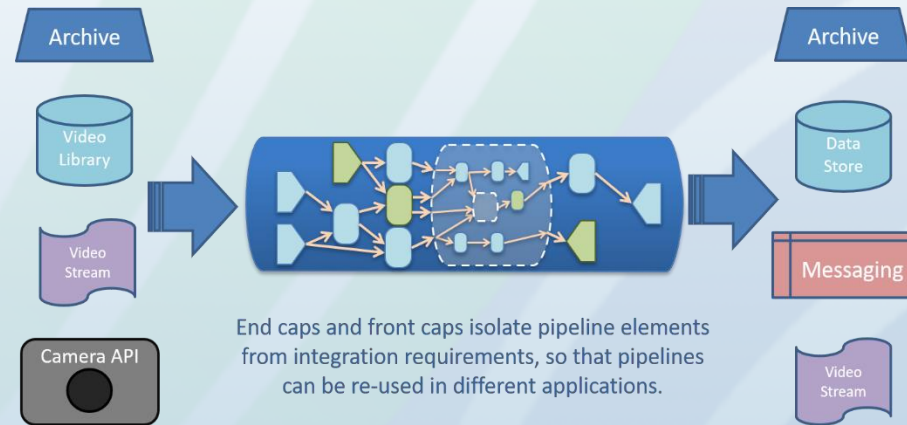
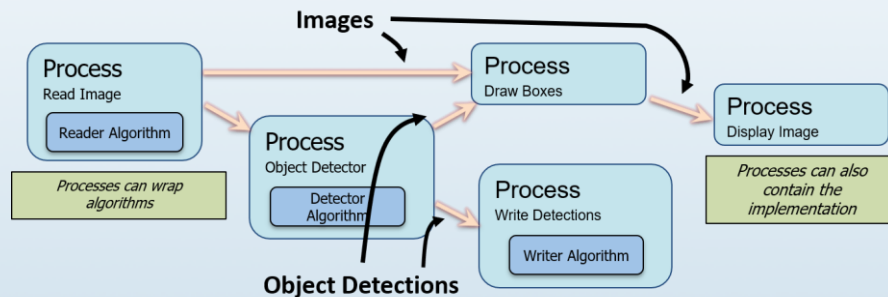


# History – Algorithm Integration Platform



Capability	Primary data source	POC	Stereo calibration	Stereo processing	Video	Color, contrast correction	scallop detection	fish detection	fish length, sizing	fish tracking	fish classification	anomaly det.	habitat classification	image segmentation
NW SC CamTrawl	Cam Trawl	Williams	yes	yes	4 Hz	no, grayscale	yes	yes	automatic	yes	yes			
ROV video fish detection and tracking	SWFSC ROV video	Cutter	no	no	30 Hz	yes	yes, DPM (UW)	yes	no	yes (UW student)	desired	desired		
ROV stereo fish measurement	SWFSC ROV GigE stereo	Cutter	yes	yes	2-4 Hz	yes	no	manual	no					
WHOI/NEFSC scallop detector	HABCAM towed rig	Dvora	yes	yes	no	yes	yes							
RPI/Kitware scallop detector	HABCAM towed rig	Hoogs	no	no	no	yes	yes							
SRI fish detection, classification, size	PI FSC MOUSS/BotCam	Ben/Mike	yes, accept cal files	yes	30 Hz	no, grayscale	yes	yes	yes	yes	yes			
SEFSC stereo proc	Drop cams from SEFSC	Thompson	yes	yes	yes		yes, basic background	manual	no	no				
Toyon SBIR I	Drop cams from SEFSC	Thompson	yes	yes	yes		yes, basic HOG	manual	yes	yes	yes			
LANL segmentation and shape analysis	HABCAM towed rig	Lakshman	no	yes	no	no	yes	yes	no	no	yes (image)	yes	yes (polygonal)	
Toyon SBIR II	Still Images AUV, drop, towed	Clarke	yes	yes	no	yes (Hanu)	yes	yes	no	yes				
WHOI/NEFSC habitat classifier	HABCAM towed rig	Dvora	yes	yes	no	yes						yes	partially	
NWFSC clustering	AUV and MOUSS	Clarke	no	no	no							yes	partially	

Green well-implemented; quantified, comparative performance assessment; ready for integration  
Yellow Existing implementation as mature research code; some performance quantification  
Red preliminary research code with ongoing work against major problems  
Gray idea or concept; no implementation



- Base classes for common operations (image filters, object detectors, trackers, ...)
- Derivation of base classes in C/C++, CUDA, Python, or Matlab
- Backend coded in C++ for efficiency, automatically multi-threaded

# Three Detection Workflows



Input Data




Running Existing  
Detectors and Trackers  
(Traditional or Deep)




New Detector  
Training via Deep  
Learning



New Classifier  
Training via IQR



New Annotations or  
Annotation Correction  
Fix Boxes, Labels, or Masks




Higher-Level Analytics

(Automated Detection Quantities, Heatmaps, Occurrence vs Time Plots)







# Types of Annotation and Detection Models




**Box-Level**



**Frame-Level**



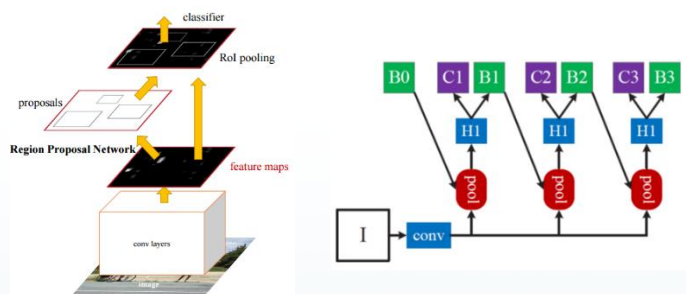
**Pixel-Level**




**Keypoints**

## Baseline Object Detectors

## Cascade Faster R-CNN [1]




## YOLOv3 and v4 [3,4]



## Cascade Mask Faster R-CNN [2]



## Scallop-TK



VIAME contains multiple baseline general purpose detectors from the larger computer vision community for wide applicability, but then specializations and other functionality added specific to domains of interest

- [1] Cai, Zhaowei, et al. "Cascade R-CNN: Delving into High Quality Object Detection." CVPR 2018.
- [2] Chen, Kai et al. "MMDetection: Open MMLab Detection Toolbox and Benchmark." arXiv preprint 2020.
- [3] Redmon, Joseph, and Ali Farhadi. "YOLOv3: An Incremental Improvement." arXiv preprint 2018.
- [4] Bochkovskiy, Alexey et al. "YOLOv4: Optimal Speed and Accuracy of Object Detection." arXiv preprint 2020.




# Automatic Parameter Optimizations

- Automatically handle LR stepping based on validation loss, running multiple hyperparameter sets, and early stopping criteria
  - NetHarn: <https://gitlab.kitware.com/computer-vision/netharn>
- Automatically choose whether to grid detectors over image
- Utilities to turn dot annotation into boxes


## Auxiliary Data Fusion

Fuse depth and motion maps into object detectors




## Ensemble Classifiers

Fuse output of detectors from different frameworks




## Extra Augmentation

Selectively augment channels differently

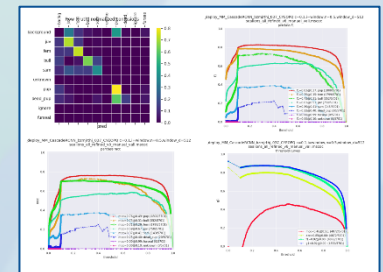


## Scoring Utilities

Evaluation via ROCs, PRCs, track metrics, ...




Detection-Pd: 0.791209  
Detection-FA: 213  
Detection-PFA: 0.515738  
Frame-NFAR: not computed  
Track-Pd: 0.791209  
Track-FA: 213  
Computed-track-PFA: 0.515738  
Track-NFAR: not computed  
Avg track (cont., purity): 1.34, 1  
Avg target (cont., purity): 1.47, 0.79  
Track-frame-precision: 0.5  
DEPTH-Hash: "a2123ede"

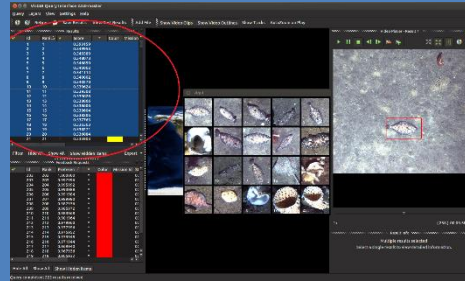


# Interactive Search and Rapid Model Generation


**Generate Image Descriptors around Detections**




**Query Image Archive using Image Exemplar**



**Create Classifier using Interactive Query Refinement**



**Generate Generic Object Detections**



**Data Storage  
via DB**

**Run Detector and Classifier on New Images**





# Interactive Search and Rapid Model Generation

VisGUI Query Interface 2.0.0-master

Query Layers View Settings Help

Refine Save Results View Best Results Add File Show Video Clips Show Video Outlines Show Tracks AutoZoom on Play

Results

✓	Id	Rank	Score	*	Color	Mission Id	Start
✓	2	2	0.601095	*	Yellow		00:16
✓	3	3	0.599851	*	Yellow		00:02
✓	4	4	0.592139	*	Yellow		00:06
✓	5	5	0.586647	*	Yellow		00:03
✓	6	6	0.583385	*	Yellow		00:15
✓	7	7	0.580398	*	Yellow		00:06
✓	8	8	0.578843	*	Yellow		00:03
✓	9	9	0.576041	*	Yellow		00:05
	10	10	0.574787	*	Yellow		00:07
	11	11	0.570809	*	Yellow		00:06
	12	12	0.568628	*	Yellow		00:24
	13	13	0.566879	*	Yellow		00:07
	14	14	0.566052	*	Yellow		00:00
	15	15	0.565552	*	Yellow		00:06
	16	16	0.565048	*	Yellow		00:07
	17	17	0.564532	*	Yellow		00:07
	18	18	0.563615	*	Yellow		00:06
	19	19	0.561714	*	Yellow		00:04
	20	20	0.561407	*	Yellow		00:06

Filter Hide All Show All Show Hidden Items Export

Feedback Requests

✓	Id	Rank	Preference	*	Color	Mission Id	Start
	202	202	1.000000	*	Blue		00:03
	203	203	0.997996	*	Blue		00:29
	204	204	0.995992	*	Blue		00:13
	205	205	0.993988	*	Blue		00:14
	206	206	0.991984	*	Blue		00:25
	207	207	0.989980	*	Blue		00:12
	208	208	0.987976	*	Blue		00:23
	209	209	0.985972	*	Blue		00:14
	210	210	0.983968	*	Blue		00:06
	211	211	0.981964	*	Blue		00:23
	212	212	0.979960	*	Blue		00:05
	213	213	0.977956	*	Blue		00:23
	214	214	0.975952	*	Blue		00:16
	215	215	0.973948	*	Blue		00:13
	216	216	0.971944	*	Blue		00:20
	217	217	0.969940	*	Blue		00:13
	218	218	0.967936	*	Blue		00:23

Hide All Show All Show Hidden Items

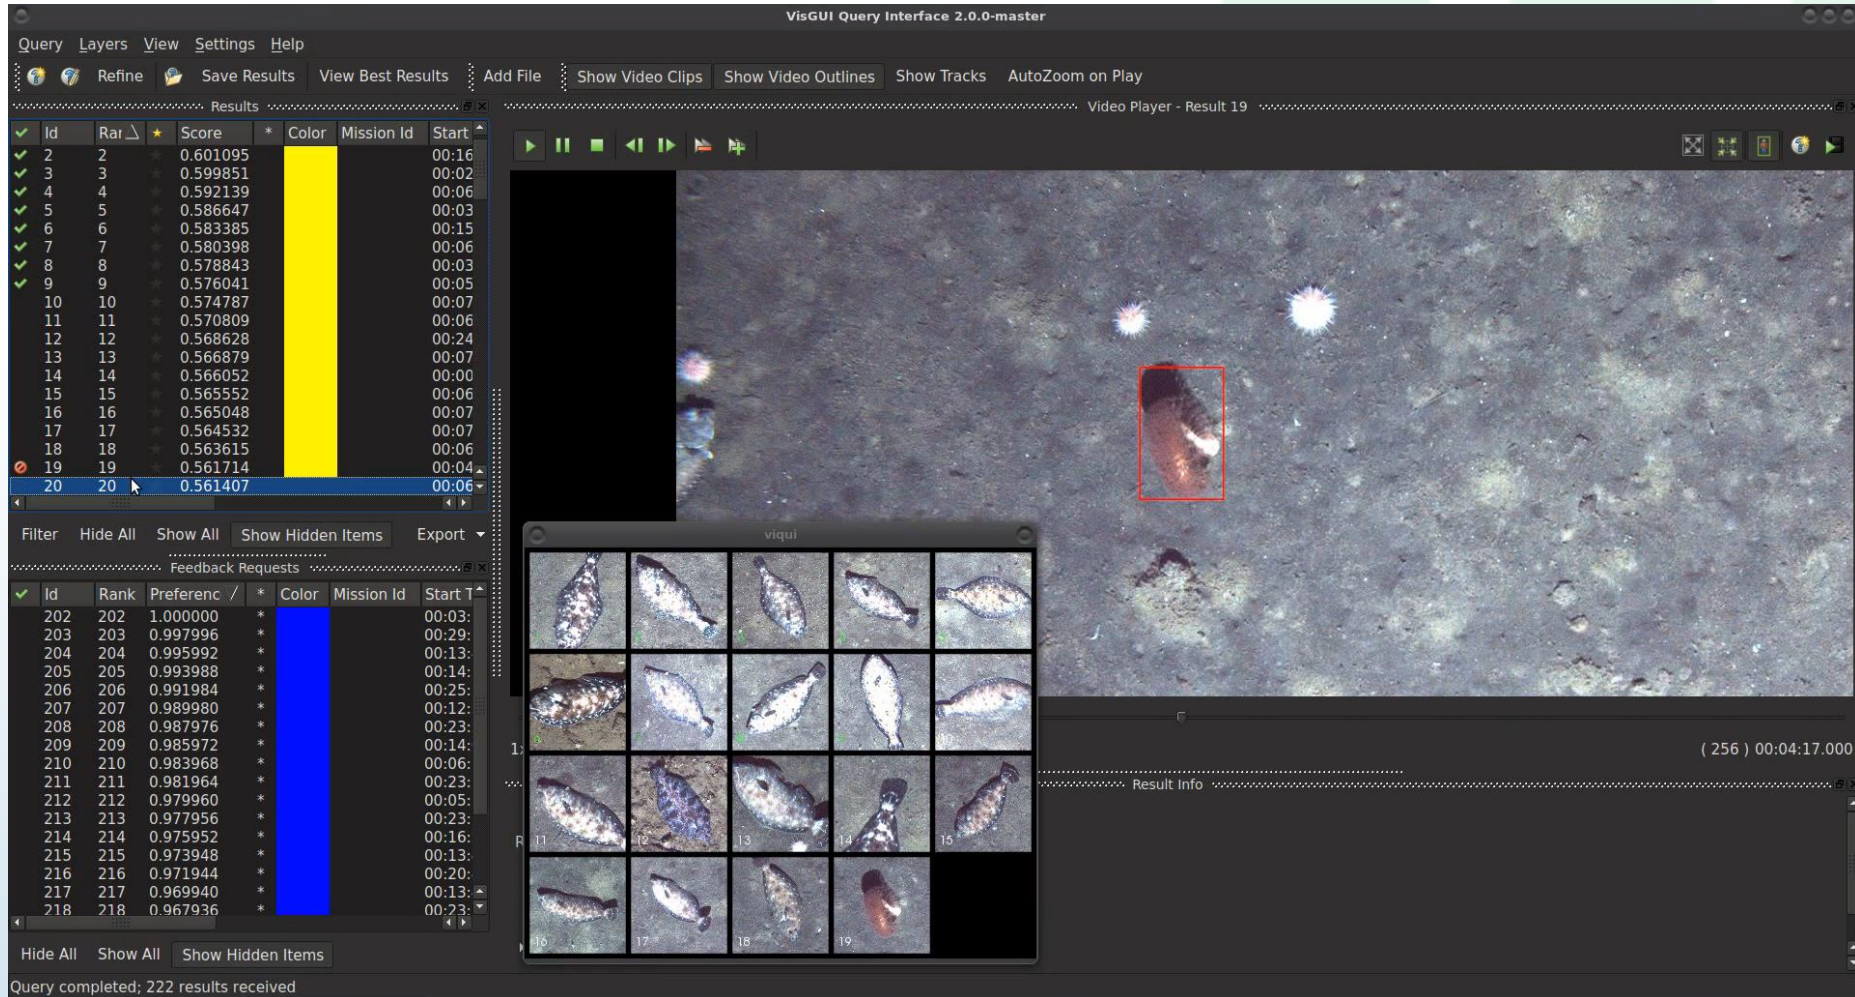
Query completed; 222 results received

Video Player - Result 19

( 256 ) 00:04:17.000

Result Info

viqi


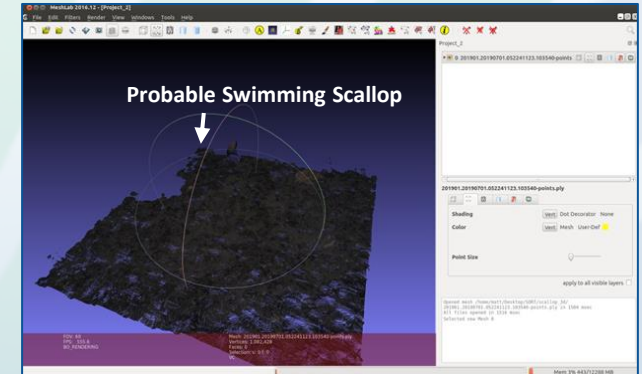
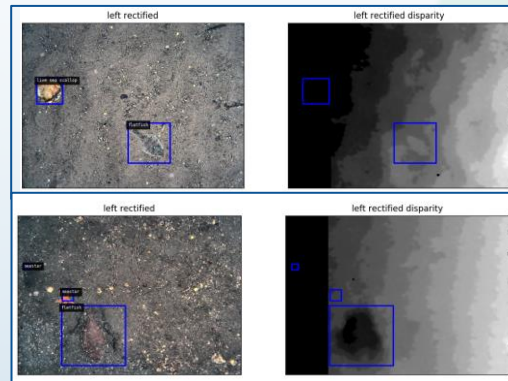



User provides  
initial image query:



User corrects system returns on  
subsequent iterations through iterative  
query refinement (IQR)



# Enhancement, Calibration and Depth Estimation



Work performed in conjunction with  
Coonamesett Farm Foundation (CFF) and NEFSC

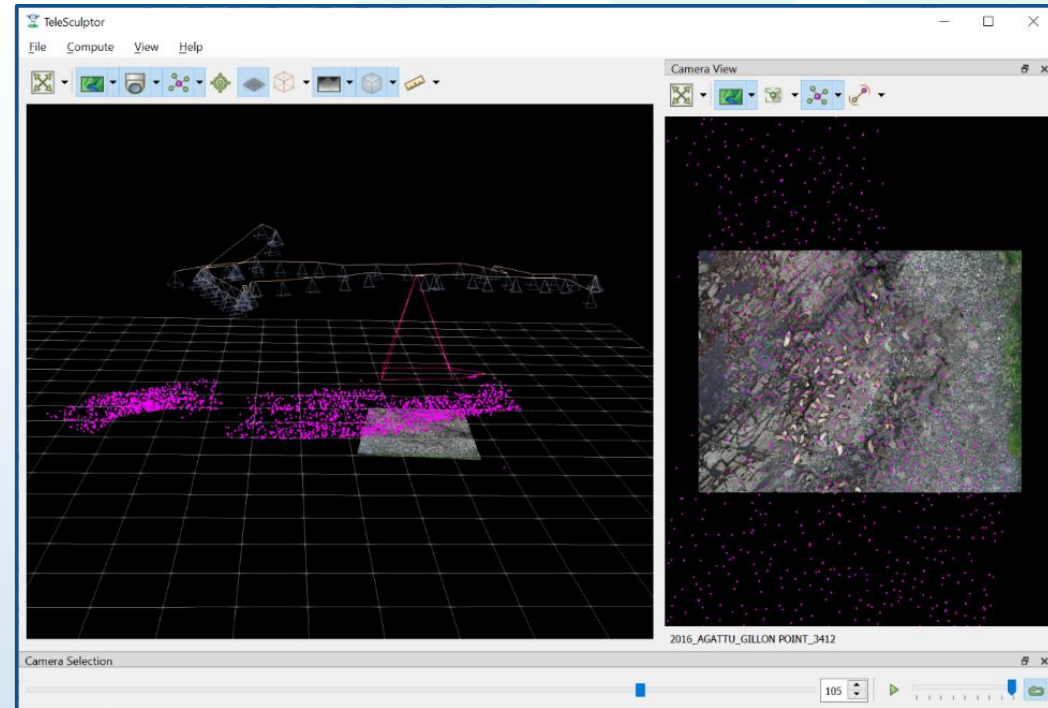



# Registration and Mosaicing



Multi-Camera and Multi-Modality Registration

Work performed in conjunction with  
NOAA AKFSC Marine Mammals Lab



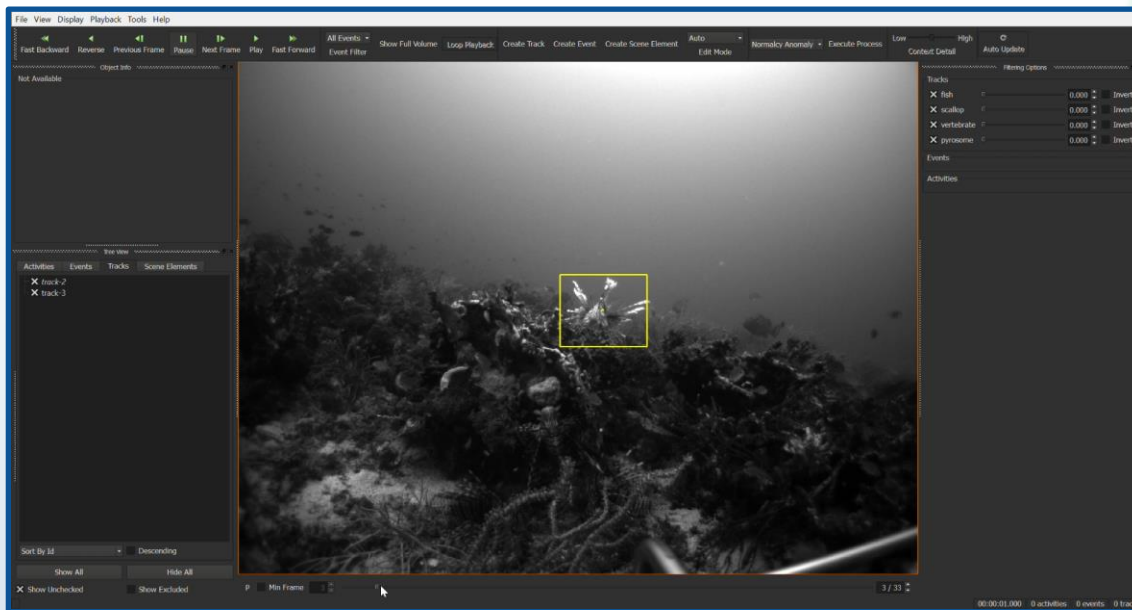



Above Source: TeleSculptor Open-Source 3D Geometry,  
<https://github.com/Kitware/TeleSculptor>  
Contact: [matt.leotta@kitware.com](mailto:matt.leotta@kitware.com)

# Object Tracking

Have integrated multiple trackers into VIAME:

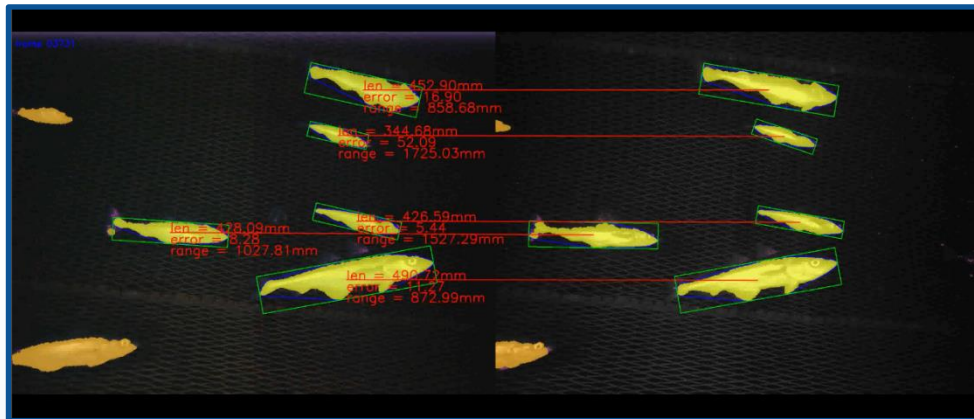
- Registration-Only Based (Aerial Pinniped)
- Deep Learning Based LSTM Tracker Detector Linkers [5]
- Non-Deep Learning Kalman Filter Detector Linkers
- Single Target Trackers (e.g. SiamRPN++ [6]) for annotation assist (below)




Code developed in conjunction with Air Force Research Laboratory (PRR# 88ABW-2019-4904, Distribution Statement A - Approved for public release: distribution unlimited).



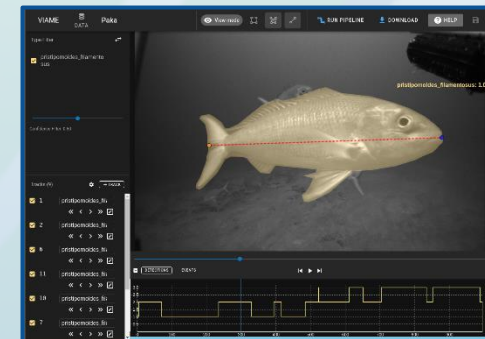
# Stereo Measurement



**Method #1** - Modeled off camtrawl process (Williams et al 2010), python port of matlab code. Use centroids of smaller sides of oriented bounding boxes as head/tail positions.




**Method #2** - Feature point detection using dedicated CNN keypoint detectors, either in the same network or separate dedicated network (e.g. heavily modified version of [7]). Ongoing: collecting additional annotations.



# Full Frame Classification

```
unannotated_frame 10000
surf_sky 0
surf_head 2
open_water 3
all_krill 0
water_open 204
krill 277
water_uso 16
water_penghead 15
surf_peng 64
surf 104
dark 55
krill_penghead 22
blasted 25
water_pengfriend 15
```



Substrate = X


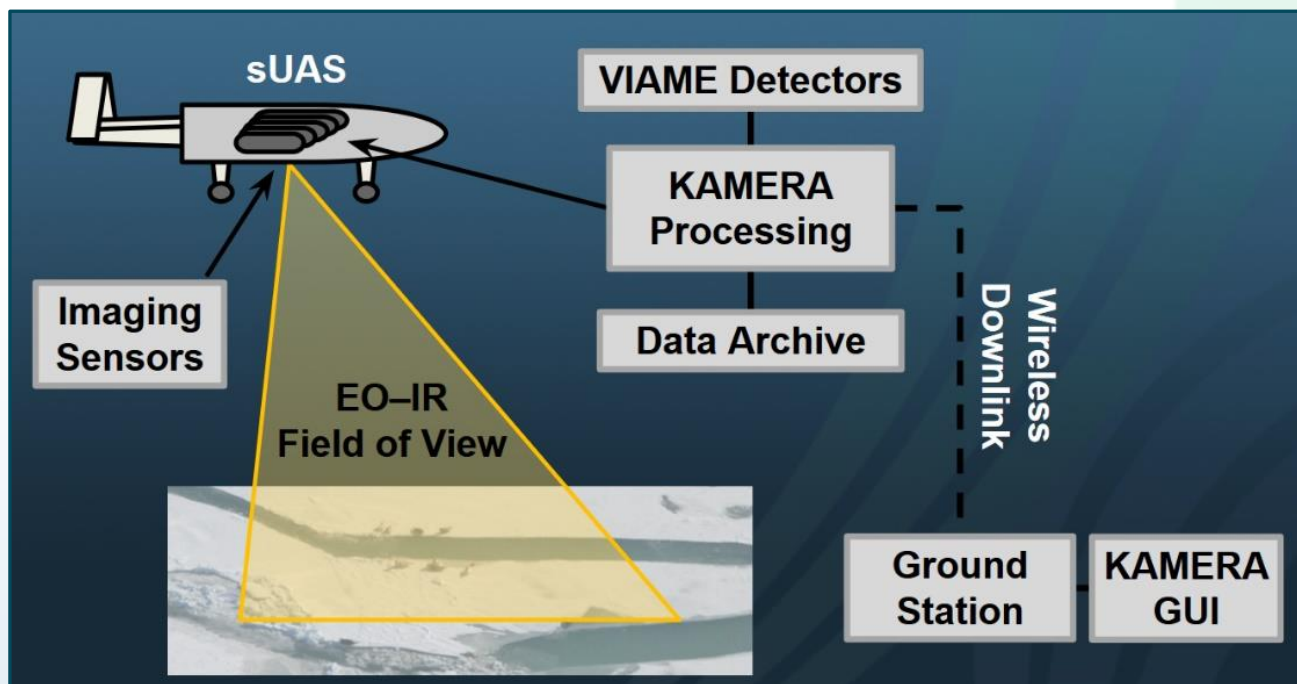
Substrate = Y

**Method #1** – Typical deep training pipeline (ResNet50 [8] – better for cases that have a lot of manual groundtruth)

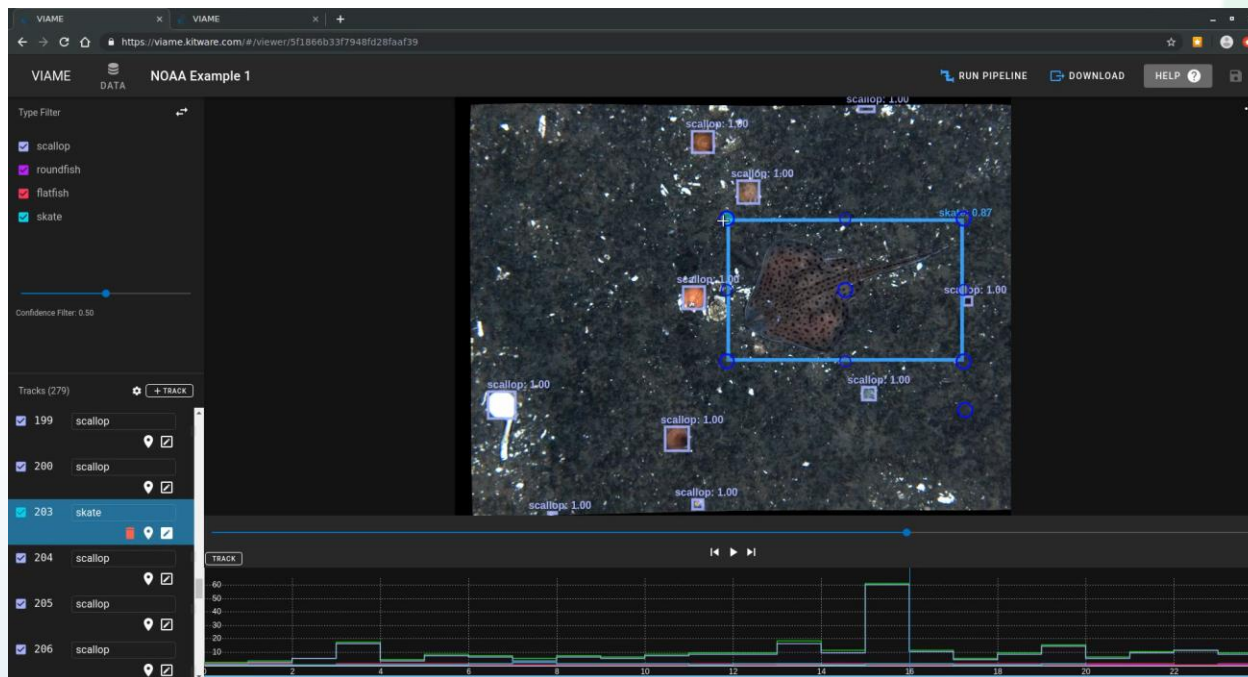
**Method #2** – SVM on fixed feature vector, similar to image search and rapid model generation pipeline (better for less training samples)



# Embedded Processing




# Graphical User Interfaces




## Web vs Desktop

- Both wrap arbitrary processing pipelines and have tracking support (split, merge)
- Web currently has better pixel classification annotation, though some in desktop
- Desktop can currently drive user initialized tracks, web can't, though adding shortly
- Desktop currently has search and rapid model generation and multi-view display, not in web
- Normalizing features across all versions in progress

## VIAME-Web Annotator



Desktop Default Annotator




Desktop Search Engine



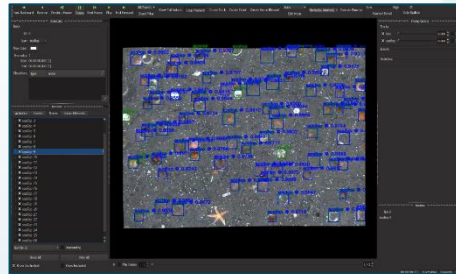
Stereo/Multi-View Annotator





# Complexity



## Desktop Applications



## Command-Line Tools and APIs:



**Covered in Tutorial on Sept 22<sup>nd</sup>, 12 pm to 2 pm EST** 

# Public VIAME Server

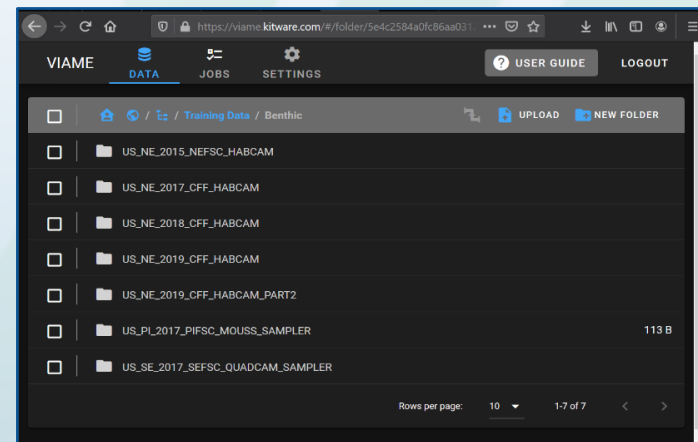
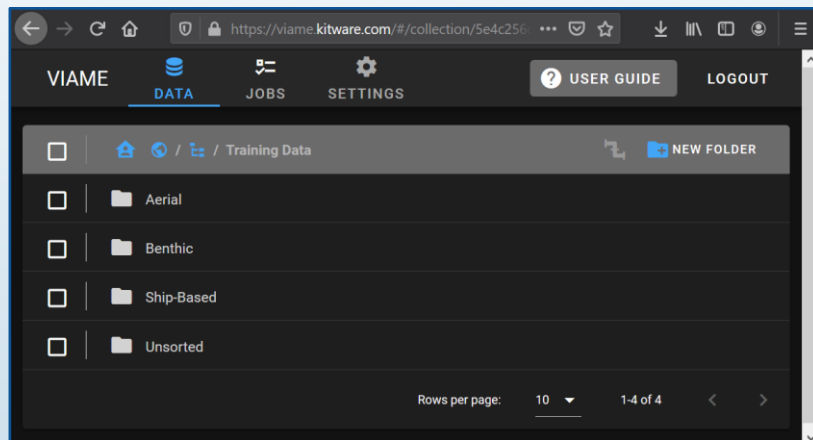
Example VIAME-Web instance: <https://viame.kitware.com>

Public data and annotation store provided by Kitware for:

- Storing and Sharing of Annotations and Imagery related to VIAME
- Performing Annotation
- Running Detectors on Data
- Limited Detector Training

**15 Tb open training data store (raided for backup), 2 GPUs**

- 3 Tb used so far, though still sorting annotations to post online
- Users can annotate data, run multiple pre-trained detectors, and (shortly) use limited training capabilities (FIFO queue on 1 GPU) for new categories

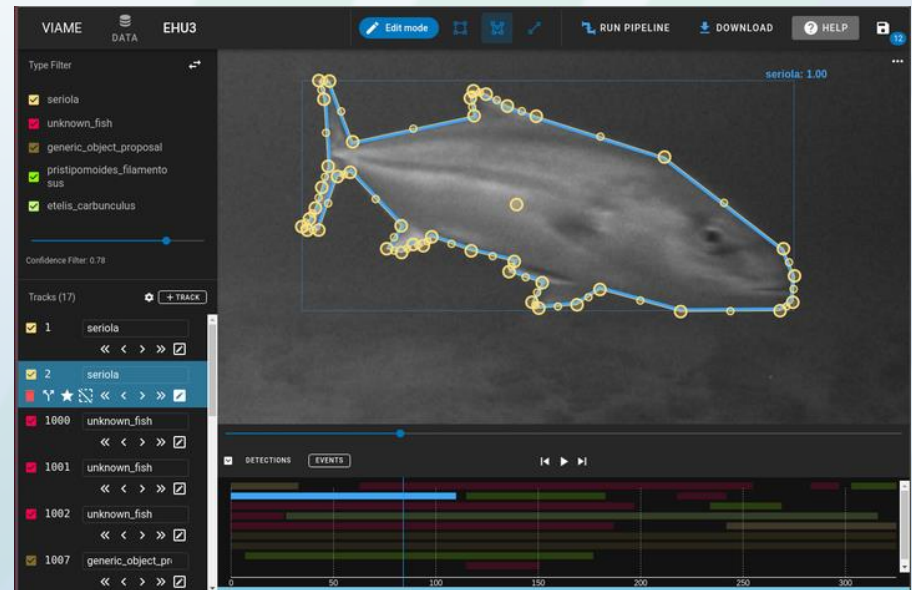


Code and docker containers: <https://github.com/VIAME/VIAME-Web>



# Current and Future Work

- Full feature support in web GUIs, improved desktop GUIs
  - Web GUI is very focused towards object detection and tracking, adding in support for other auxiliary features, such as mosaic generation, image enhancement, etc...
    - Make cross-use of other projects at Kitware performing similar work (IARPA, DARPA, DoD)
  - Pixel-classification utility pipelines (e.g. boxes to masks)
- Additional algorithm specializations
- More documentation
- Behavior and event detection
- Acoustic data processing
- Electronic monitoring




# Conclusions




- VIAME is a do-it-yourself (DIY) AI toolkit which can be applied to multiple types of imagery or video
- Can be run by people with no programming or machine learning background in both web and desktop interfaces
- Released as fully open-source with a permissive license
- Specializations to maritime processing such as motion fusion, stereo measurement, image enhancement, and object tracking which other software (e.g. Amazon SageMaker) lack

**Object Detection**




**Object Tracking**




**Image Enhancement**




**Image Registration and Mosaicing**





**Video Search and Rapid Model Generation**



**Stereo Measurement**



**Calibration, 3D and Altitude Estimation**





## Special Thanks to:

- NOAA AIASI, NOAA AKFSC MML and CFF for funding VIAME
- Image annotators and testers across various organizations

**See tutorial (Sept 22<sup>nd</sup>) and  
<https://viametoolkit.org/> for more information**

### Tutorial Agenda

12:00 to 12:15 pm	–	VIAME Overview Desktop Installation Web vs Desktop vs APIs
12:15 to 12:45 pm	–	VIAME Web Edition Types of Annotation Model Training
12:45 to 1:45 pm	–	VIAME Desktop Edition Core Functionalities Auxiliary Features
1:45 to 2:00 pm	–	APIs and Configs [Advanced]
2:00 to 2:30 pm	–	Open Discussion